ESiWACE WP3

Final Review

M. Carter[1], K. Serradell[2], R. Budich[3]

S. Kosukhin[3], D. Matthews[1], G. Greed[1], S. Fiore[4]

Presenters: Kim | Glenn

Addresses: 1) MO, 2) BSC, 3) MPI-M, 4) CMCC

Corresponding Author: reinhard.budich@mpimet.mpg.de

Topics

- Results achieved
- Technical highlights
  - Spack: Status of the tool
  - Cylc: The experimentation driver
  - On the way to exascale
- Challenges
  - Continuation of efforts and maintenance
  - Adoption into experiments
- Conclusions & Outlook
  - Link with ESiWACE2

WP3 Results achieved

- By adapting and employing the SW package **SPACK**, Substantial improvement of the software, computing and data-handling infrastructure for ESM scientists from the applications through the software stack to the hardware

- **Cylc** has been further developed to support the ever-increasing scale and complexity of ESMs for research, production and real-time operations in weather and climate

- New **analytics** workflows capabilities have been properly implemented to meet additional users requirements

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

SPACK

Workflow

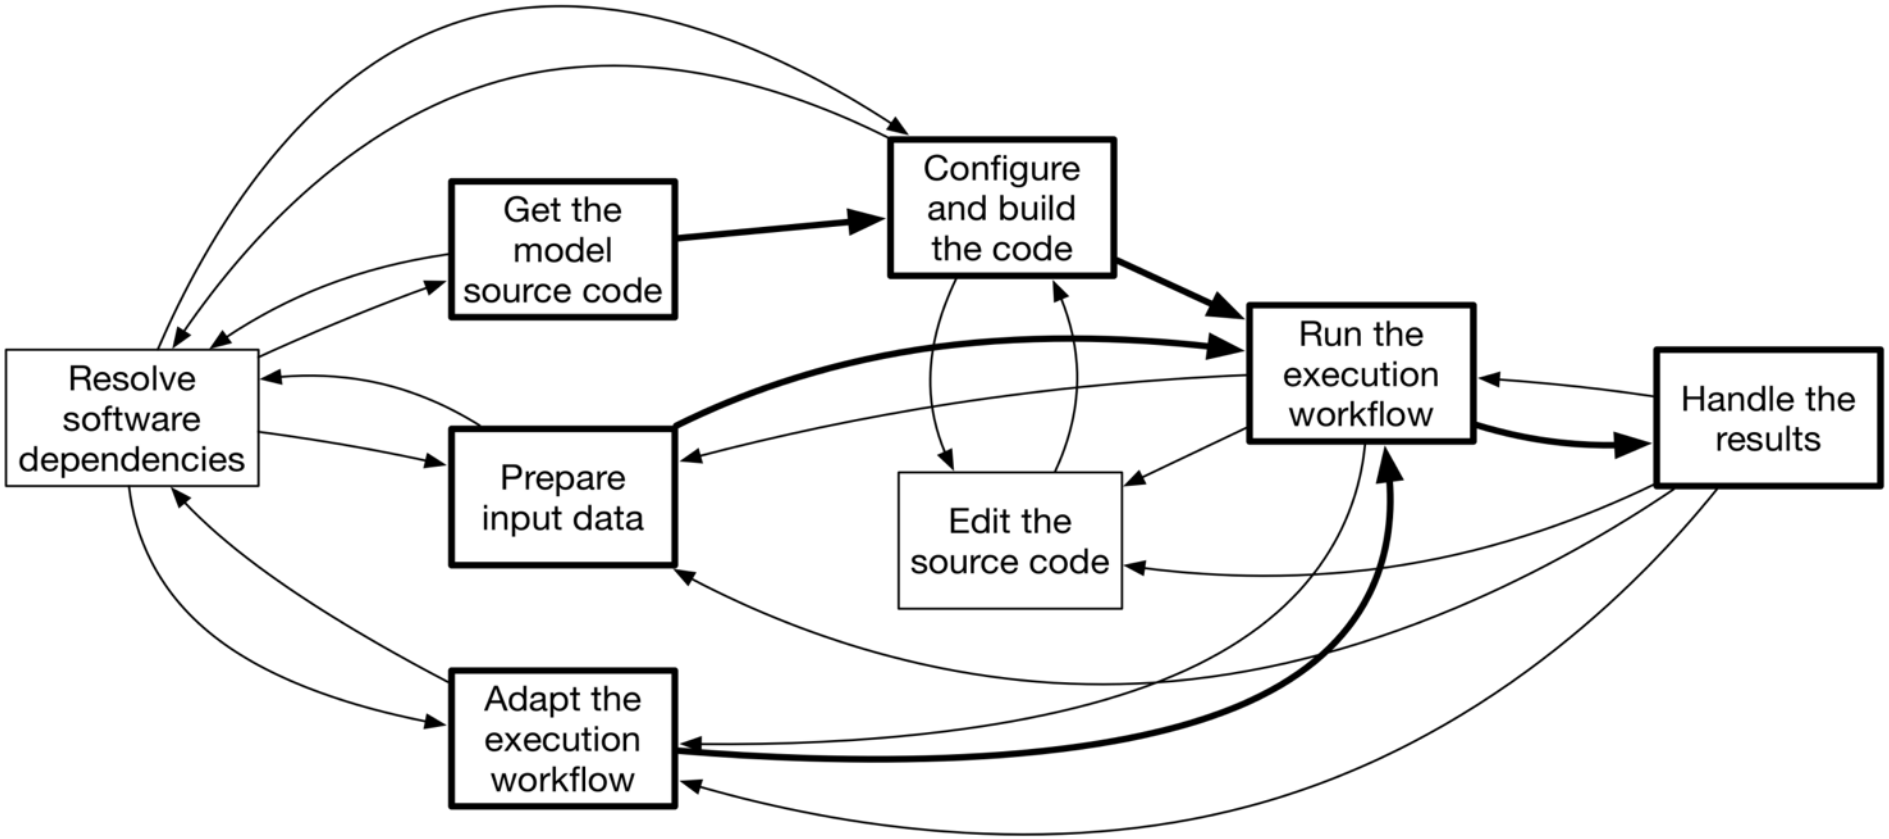

Courtesy: Sergey Kosukhin

Problem

Fuzzy questions that need to be answered before an experiment can be started

Different environments

Single machine or HPC site?
What software is already there?
What privileges do I have?

Different requirements

Which versions?
Which features to enable?
Which API implementation?

Transition between areas of responsibility

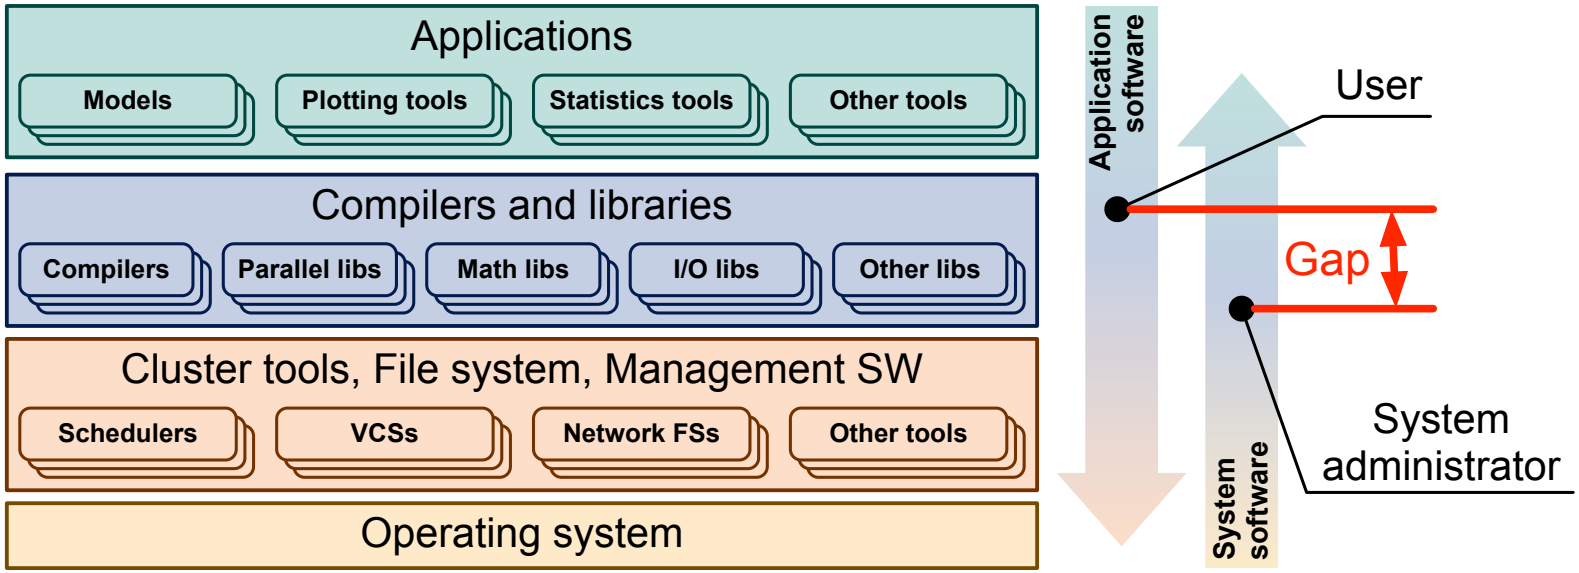

Courtesy: Sergey Kosukhin

Spack: Target Audience

- **For users:**
  - Get the productive software environment for their workflows with minimal lag
- **For system administrators:**
  - Understand the requirements of the scientific software
  - Automatization of the maintenance of the software stack
  - Both with the support of a large (and growing) community
- **For developers:**
  - Simplify the maintenance of their code
    - no need to provide bundled libraries and complicate installation scripts
  - Ensure portability of their code
  - A tool for testing their software with different compilers, libraries, etc.

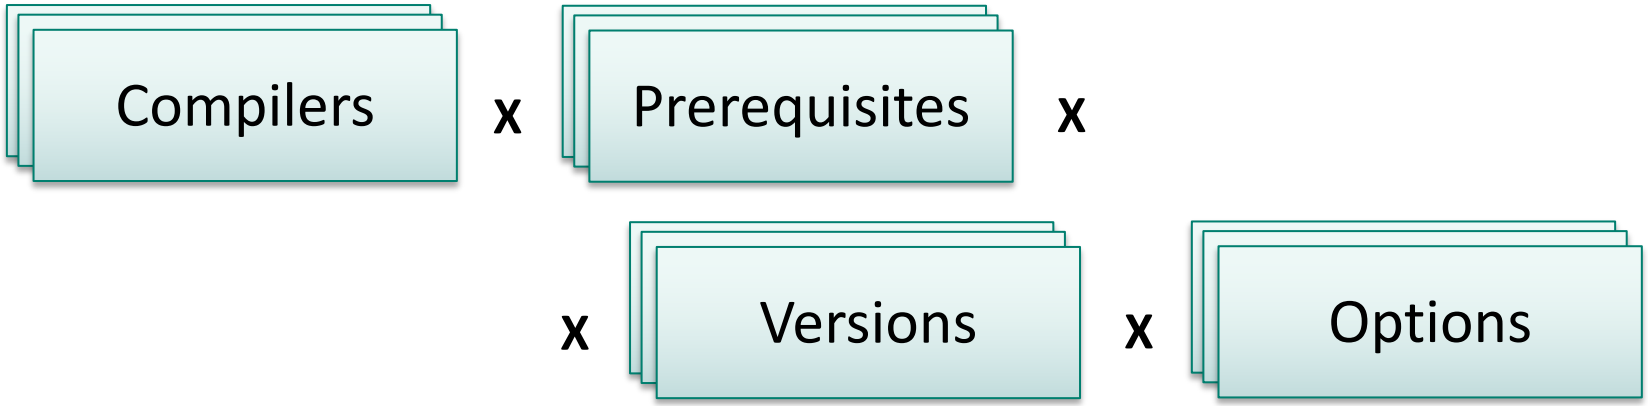

Courtesy: Sergey Kosukhin

Software stack

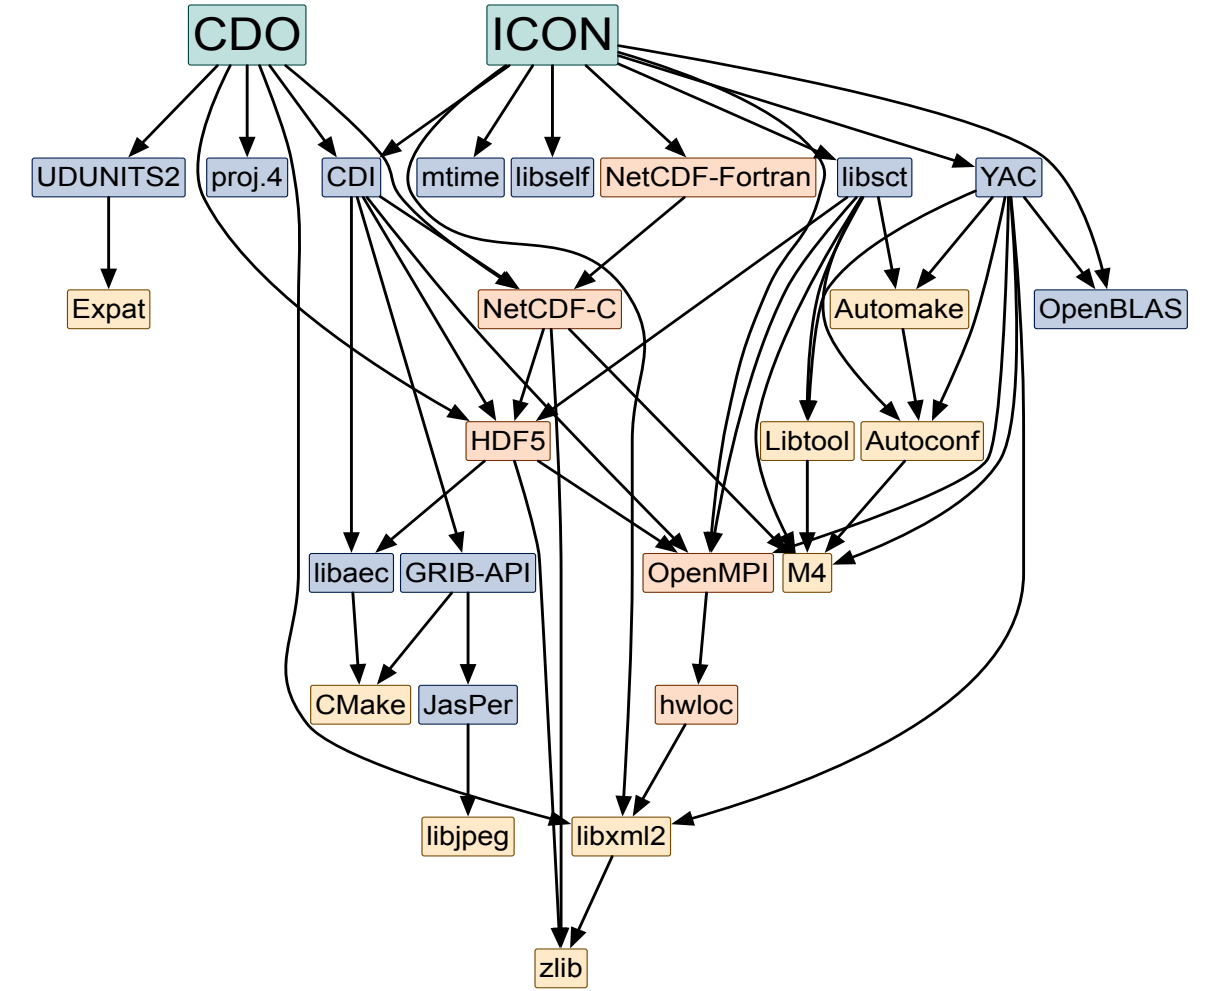

spack

Courtesy: Sergey Kosukhin

Spack for ESM

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Commonly used packages

CDO, CMOR, GRIB-API, ecCodes, Emoslib, Magics, NCL, libAEC, Extrae, Paraver, OpenBLAS, LAPACK, NetCDF, HDF5, Python (with modules), NCO, and more… (~60 accepted pull requests)

Tested machines:

- Mistral (DKRZ)
- MareNostrum (BSC)
- Piz Daint (CSCS)
- Marconi (CINECA)
- ARCHER (NCAS)
- XCE (DWD)
- Altamira (IFCA)
- Nimbus (AEMET)
- Various UNIX/Linux workstations

Courtesy: Sergey Kosukhin

Conclusions

- Users will be able to easily customise the software environment on their own, thus being more productive and reducing the workload put on system administrators.
- Spack allows for formal description of the software stack available on a supercomputer, which simplifies identification of the list of missing software dependencies of a modelling workflow.
- Spack can also help system administrators to easily test various usage scenarios of the basic elements of the software environments, such as compiler toolchains and MPI libraries.
- A good way to document the installation procedures and maintain a collection of patches (e.g. for libtool).
- Software quality was increased by the activity

Courtesy: Sergey Kosukhin


esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER
AND CLIMATE IN EUROPE

CYLC

Cylc engine

- Cylc is an Open Source workflow engine that orchestrates cycling systems very efficiently

- It is a generalized tool suitable for many environments but most often used in Weather and Climate

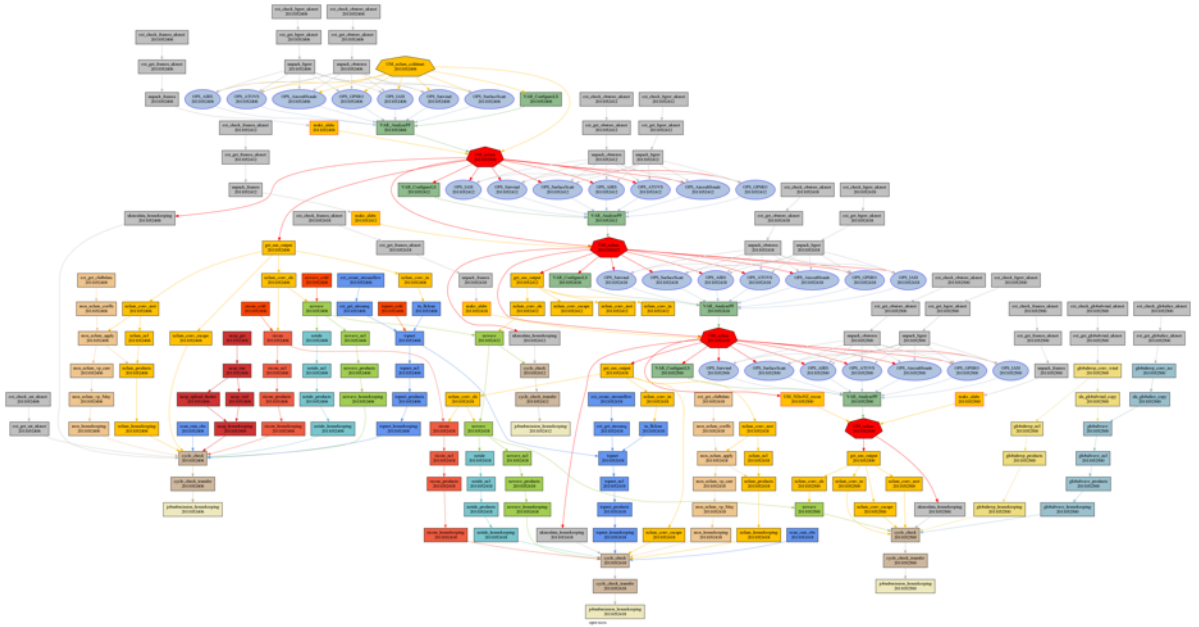

Cylc achievements

- Very active development
  - Significant improvements to performance, efficiency, robustness, testing & documentation
  - ~850 pull requests (fully tested & reviewed changes)
  - Part funded by ESiWACE
- Relatively little demand for support – good documentation
- ENES Workshops on Workflows
  - Lisbon, September 2016 (including Cylc tutorial)
  - Brussels, September 2018 (including Cylc consultancy session)
- Notable uptake
  - Adopted by NCAR for running CESM simulations
  - Altair providing support for Cylc to the Bureau of Meteorology, Australia

Courtesy: Dave Matthews

Developments in final period include …

- Scalability enhancements (needed now and for Exascale)
  - Optimized client-server interactions
  - Reduced memory usage **to half**
  - Quicker validation at start-up for tasks in suites is now **6x faster**
  - Improved sub-process management, less interruptions
  - Suites distributed across hosts and auto-migrated for server maintenance.
- Usability enhancements
  - Improved documentation to support adoption
  - A web-based job log viewer & simpler service files for easier monitoring of suites
  - Complex suites easier with parameterised tasks

Cylc Sustainability

- 20+ sites demand it!
- Open user & developer forum https://cylc.discourse.group/
- GitHub used for development https://cylc.github.io/

- Core contributors at NIWA (New Zealand) and the Met Office
  - But staff losses at MetO need to be replaced.
  - Funding secure, recruitment ongoing.

- Enhanced by EU funding but not reliant on it
  - By far the most resources funded by NIWA and MetO, but ….
  - **IS-ENES3**
    - WP6 - VA1 (11PM) – General support services
    - WP8 - JRA1 (24PM) – Community driven development
  - **ESiWACE2**
    - WP3 (6PM) - Cylc support services (via service call)
    - WP4 (12PM) - Enhance Cylc to deal with data dependencies and lifecycle information

Lessons learned

- Adoption is the main challenge
  - It is about people and culture not the technology;
  - Success is most likely when you:
    - Have an active advocate locally prepared to help people;
    - You select the right project with appropriate and motivated technical team members;
    - You have institution level backing, prepared to find appropriate resource levels
  - Much adoption has needed relatively little support
    - Good documentation supports this
  - The authors believe that as workflows become more complex as we move to Exascale, workflow tools will become more necessary.
    - But the initial investment is hard compared with incremental development of traditional approaches

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

THE ADOPTION CHALLENGE

From Autosubmit to Cylc

- BSC is developing our own workflow manager tool since 2010
- Although used in production, it has been hard to find external users (sustainability of tools)
- In the meantime, Cylc tool is being adopted as the workflow manager tool by the community
- First task: identifying pros and cons (GUI, families, improved log management, Rosie to manage suites, …)

Courtesy: Kim Serradell

Cylc and EC-Earth

- Climate Prediction
  - EC-Earth / NEMO / OpenIFS
    - Historical runs
    - Seasonal / decadal runs
    - Model development
  - Portable experiments including
    - Compilation
    - Initialization
    - Simulation
    - Post processing
    - Automated transfers
    - Data assimilation

- Rose suites have been created for running EC-Earth historical and multi-member ensemble climate simulations

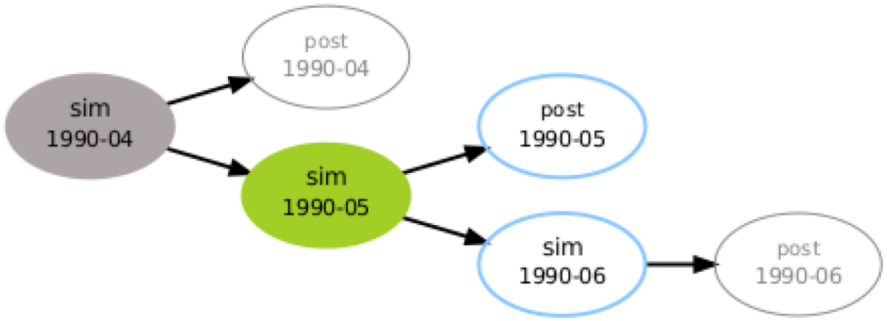

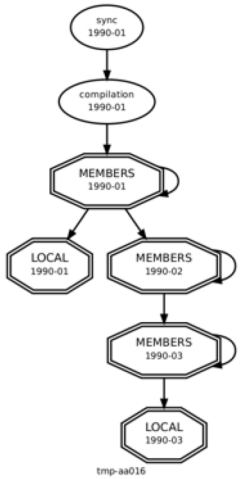

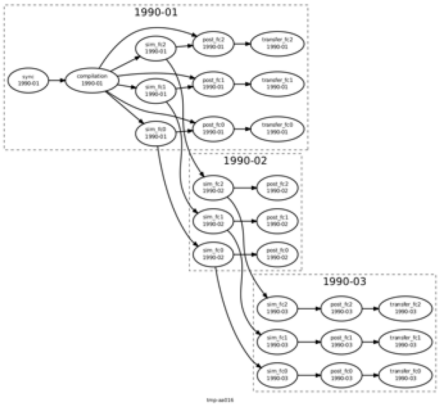

Courtesy: Kim Serradell

Lessons learned

- It took long time to take a decision and write a strategy
- Some tweaks/workarounds needed and some features are needed
- The adoption of a new production tool is a long and difficult process
- Never underestimate the *“human factor”*
- We need to find the resources not only to sustain the development of codes but we also need to ease adoption.

Courtesy: Kim Serradell

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

ANALYTICS WORKFLOWS

Two-level workflow strategy for a data analytics use case

- Use case on Fire Danger Prevention
  - Based on the operational service running at CMCC for Fire Danger Prevention
  - Serving regional stakeholders (Apulia Regional Civil Protection Department)
  - Joining weather forecast (WRF model - 5.5x5.5 km resolution) and data analytics (to compute fire danger indices/maps over the Apulia Region)
  - Successfully 'ported' on Cylc (previously based on bash scripts)
  - Implementing a two-level workflow strategy based on Cylc and Ophidia

- Results: better orchestration layer
  - Implementation of a two-level workflow strategy joining:
    - Cylc at level1 (coarse grain, general purpose) and
    - Ophidia at level2 (fine grain, on the analytics side)
      - Cylc triggers - with a single level1-task - an Ophidia workflow which consists of tens of level2-analytics tasks

- Knowledge transfer and further exploitation
  - The successful use case has been demonstrated internally at CMCC as a pilot
  - As a result, other operational services workflows have been ported to Cylc
  - The knowledge transfer to the Modelling Divisions has increased the Cylc exploitation

Workflow: cylc + Ophidia

Courtesy: Sandro Fiore

Analytics workflow support

- Improvement of the Ophidia workflow capabilities
- Driven by users' requirements (climate indicators computation and multi-model analysis)
- Improvements relate to:
  - loops (both iterative and parallel), conditional constructs, interactive support and cross-workflows synchronisation.
  - token management check for long-running workflows
  - new filters for HTC-based operators (parameter sweep jobs) negation support in filter for massive operations
  - Workflow checker improvement
  - Overall, **82 commits** have been performed on the Ophidia server module which is responsible for the workflow engine aspects), **6 issues** have been closed and **10 pull requests** have been merged
  - Training session on Ophidia during the Workshop in Brussels, September 2018

NAME	DESCRIPTION
OPH_ELSE	Start the last sub-block of a selection block "if".
OPH_ELSEIF	Start a new sub-block of a selection block "if".
OPH_ENDFOR	Close a loop "for".
OPH_ENDIF	Close a selection block "if".
OPH_FOR	Implement a loop "for".
OPH_IF	Open a "if" selection block.
OPH_INPUT	It sends commands or data to an interactive task.
OPH_SET	Set a parameter in the workflow environment.
OPH_WAIT	Wait until an event occurs.

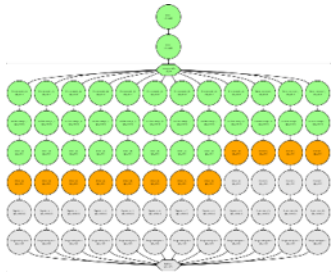

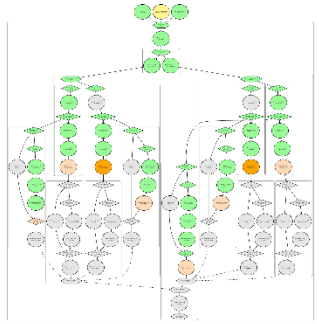

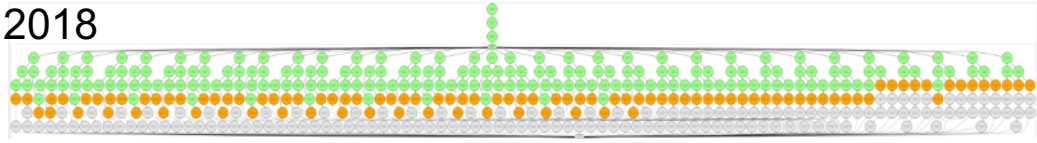

Analytics Workflow: Ophidia

Courtesy: Sandro Fiore

Multi-model analysis workflow Architectural patterns (I)

- Two architectural patterns have been defined for task management to implement a multi-model analysis experiment as a workflow
  - test case: multi-model analysis on precipitation trend analysis
  - reference context in terms of experiment: CMIP5
  - reference infrastructure: ESGF
  - *distributed* vs *centralized* architectural approaches
  - test case: centralized approach (Climate AnalyticsHub, *Fiore et al, BigData2018*)

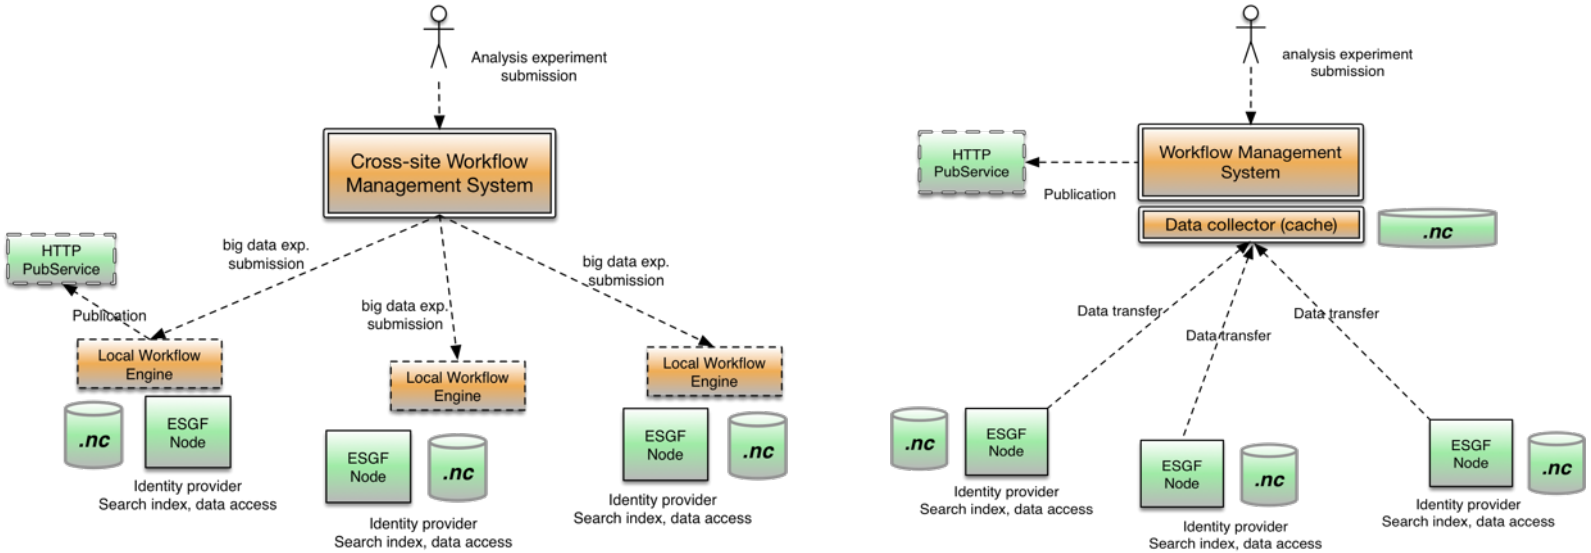

Results reported in "*D3.10 on Scheduler development and support activities*"

Courtesy: Sandro Fiore

Multi-model analysis workflow Architectural patterns (II)

- Centralized approach implemented at CMCC and based on analytics workflow
- Initially deployed on a dedicated analytics cluster
  - 100 cores, 1.5TB RAM, 50TB storage
- Ported on the Athena HPC production cluster
  - 482 nodes, 16 cores/node, 64GB RAM/node)
  - Integration aspects: scheduling (LSF) and parallel file system (GPFS) at the cluster level
- Application test case:
  - multi-model precipitation trend analysis
  - Executed and validated on 15 models from CMIP5
  - Analytics workflow based on Ophidia
- Results presented at the HPC Summit Week 2019 (HPDA Workshop)
- Results reported in "*D3.10 on Scheduler development and support activities*"

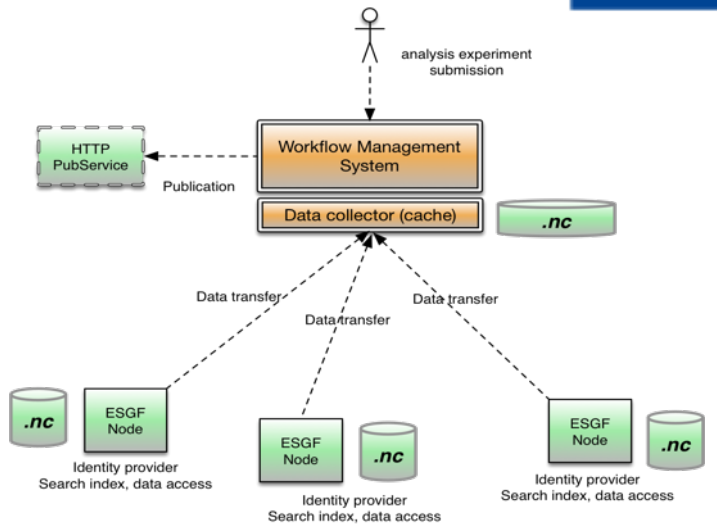

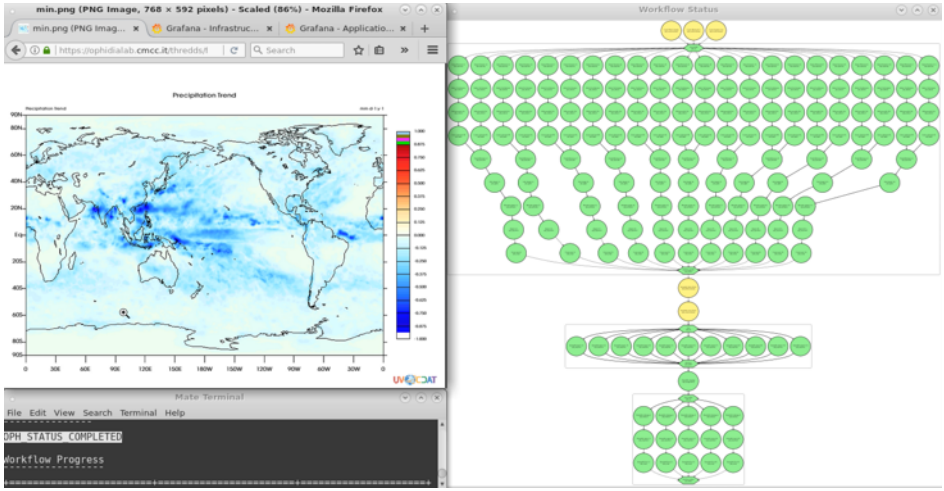

S. Fiore, D. Elia, C. Palazzo, A. D'Anca, F. Antonio, D. N. Williams, I. T. Foster, G. Aloisio: Towards an Open (Data) Science Analytics-Hub for Reproducible Multi-Model Climate Analysis at Scale. BigData 2018: 3226-3234

Courtesy: Sandro Fiore

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

CONCLUSIONS

Conclusions

- Good tools have been developed based on open source initiatives and input from outside the project
- Adoption of the tools is limited because no one use case makes it essential
- Organisational and human factor rather than technical issues dominate
- Hard to allocate resources to develop and spread the usage of the tool at the same time.
- Is this a topic for sustainability coordinated across IS-ENES and ESiWACE activities?

Link with ESiWACE2

- Significant work done to prepare software, computing and data-handling infrastructure to Exascale era
- Tools WP3 will be used and supported in ESiWACE2
  - Cylc as workflow manager
  - SPACK to deploy demonstrators (ICON, openIFS, …)
- From the usability point of view, ESiWACE has become a testbed to validate/discard technical choices